



73
B1
4.7.02

PATENT
P51380US

IN THE U.S. PATENT AND TRADEMARK OFFICE

Applicant: Adam Peter GAWNE-CAIN Conf.: 7374
Appl. No.: 09/987,592 Group: 2171
Filed: November 15, 2001 Examiner: UNASSIGNED
For: DATABASE MANAGEMENT SYSTEMS

CLAIM TO PRIORITY

RECEIVED
MAR 21 2002
Technology Center 2100

Assistant Commissioner for Patents
Washington, DC 20231

Date: March 14, 2002

Sir:

Applicant(s) herewith claim(s) the benefit of the priority filing date of the following application(s) for the above-entitled U.S. application under the provisions of 35 U.S.C. § 119 and 37 C.F.R. § 1.55:

<u>Country</u>	<u>Application No.</u>	<u>Filed</u>
UNITED KINGDOM	0028311.9	November 21, 2000

Certified copy(ies) of the above-noted application(s) is(are) attached hereto.

Respectfully submitted,

YOUNG & THOMPSON

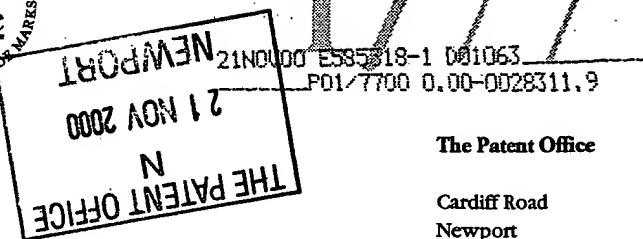
Benoit Castel

Benoit Castel, Reg. No. 35,041

745 South 23rd Street
Arlington, VA 22202
Telephone (703) 521-2297

BC/mdp

Attachment(s): 1 Certified Copy(ies)

**Request for grant of a patent**

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

 Cardiff Road
 Newport
 South Wales
 NP10 8QQ

1. Your reference

RPH.P51380GB

2. Patent application number

(The Patent Office will fill in this part)

0028311.9

3. Full name, address and postcode of the or of each applicant (underline all surnames)

 Gawne Cain Research Ltd.
 The Cottage
 5 Beauchamp Lane
 Oxford
 OX4 3LF

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

United Kingdom

8025645001

4. Title of the invention

"Database Management Systems"

5. Name of your agent (if you have one)

Marks & Clerk

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

 4220 Nash Court
 Oxford Business Park South
 Oxford
 OX4 2RU

Patents ADP number (if you know it)

727 1125 001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

 Country Priority application number
 (if you know it) Date of filing
 (day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

 Number of earlier application Date of filing
 (day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

Yes

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form	0
Description	12
Claim(s)	4
Abstract	1
Drawing(s)	4 4

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

Any other documents
(please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature *Marks & Clerk* Date

Marks & Clerk

20 November 2000

12. Name and daytime telephone number of person to contact in the United Kingdom

Richard Harding - 01865 397900

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.



INVESTOR IN PEOPLE

Docket No.: P51380US
Applicant: Adam Peter GAWNE-CAIN
Appl. No.: 09/987,592
Filed: November 15, 2001
For: DATABASE MANAGEMENT SYSTEMS
Young & Thompson
(703) 521-2297



The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

RECEIVED

MAR 21 2002

Technology Center 2100

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed:

Dated:

8 November 2001

~~DUPLICATE~~**"Database Management Systems"**

This invention relates to a database management system, and is concerned more
5 particularly with such a database management system for maintaining chunks of data
indicative of the states of the database both before and after a transaction modifying the
state of the database.

Generally, the method used by a conventional database management system
10 (DBMS) to maintain a computer database over time involves adding, modifying and
deleting data records. Since the data records are modified and deleted, the previous
states of a mature database cannot be directly revealed.

In certain known systems the accessing of old versions of relational databases is
15 possible by recording all transaction instructions in accompanying log files, along with
copies of the database file at certain checkpoints. However, this method is
unsatisfactory, since the log files must be laboriously replayed from the previous
checkpoint copy, and the checkpoint copies and log files take up a lot of space, since
they contain so much logically redundant data.

20

It is an object of the present invention to provide a structure for a DBMS to
enable a database to be maintained over time in such a manner as to allow a large
number of previous states of the database to be directly revealed in a particularly
straightforward manner.

25

According to the present invention there is provided a database management
system for maintaining chunks of data indicative of the states of the database both
before and after a transaction modifying the state of the database, the system
comprising:

30 (a) memory means for holding permanent records of the state of the database before
the database-modifying transaction and the state of the database after the database-
modifying transaction;

(b) relation determination means for relating at least one parent data item in the record of each database state to at least one dependent data item in the record of the same database state;

(c) root determination means for determining the position of a root data item in the record of each database state to which other data items in that record are related; and

5 (d) state determination means for determining the state of the database after the database-modifying transaction by relating the root data item corresponding to that database state to both at least one data item in the record of that database state and at least one data item in the record of the state of the database before the data-modifying

10 transaction as determined by the previous state location means.

Such a DBMS enables a computer database to be maintained over time in such a manner as to allow all previous states of the database to be directly accessed and without requiring the storage of large amounts of redundant data.

15

The structure of such a DBMS may be likened, by analogy, to the way in which many types of tree grow by laying down a new ring of wood each year. This structure means that all previous states of a mature tree can be directly revealed by stripping away the tree's outer rings. In an analogous way, the present invention relates to a

20 DBMS by which a computer database can be flexibly maintained over time, preferably by exclusively adding data items to the database, and preferably without requiring any existing data items to be modified or deleted.

Preferably, during each database-modifying transaction, a chunk of new data items are appended to the end of the growing database file. Previous data items in previous chunks in the file will not need to be modified. In this way, all previous states of the database can be directly revealed by the DBMS by only considering a certain number of the initial chunks, whilst ignoring the remainder of the appended chunks. For example, if the database file consists of seven chunks, corresponding to seven

25 transactions, then the state of the database after the third transaction can be revealed by the DBMS examining the first three chunks, and temporarily ignoring the last four chunks.

Preferably, the records within the chunks are structured according to the following principles:

- Chunks of data items are appended to the database file as transactions are made
- 5 • Chunks are never edited once they have been appended to the file
- Chunks are on average small compared to the total size of the previous database file.
- Chunks refer to data items in previous chunks where those data items have not been logically deleted or modified by the transaction.
- 10 • Chunks mask out data items in previous chunks where those data items have been deleted or modified by the transaction.
- Chunks do not reference data items in subsequent chunks.

These principles are preferably achieved in the following manner:

- 15 • Each data item has a position in a chunk in the file.
- Data items may be parent data items.
- Each parent data item contains data indicating the position of its dependent data items.
- Each transaction chunk contains the position of a root data item.
- 20 • The DBMS can traverse a network of data items by tracking dependent data items from a root data item, and then recursively tracking more dependent data items from the data items already visited.
- As the DBMS traverses this network of data items, it can suitably include means for recording the position of the parent data item of each visited dependent data item.
- 25 • Thus dependent data items in chunks in the database file do not store the positions of their parent data items.
- Depending on the root data item used to initiate a traversal, the DBMS may reveal a different network of data items.

30 Generally, in computer science, a database can be said to have a "physical" state, and a "logical" state. The physical state is the arrangement of data stored in a suitable

computer data storage device, such as a magnetic disk or a CD-ROM. A logical state is a conceptual view of the information embodied by these data. In general, a DBMS is a computer program that can manipulate the physical state while presenting a possibly different logical state to other computer programs.

5

The DBMS in accordance with the present invention can present a modified logical database state, wherein data items have been added, edited or deleted, by adding data items to the physical state of the database, without the need to edit or delete existing data items in the physical state.

10

Preferably, during a database-modifying transaction, the DBMS will prepare a chunk of data items to be appended to the database file in the following manner:

15

- New dependent data items can be inserted into the new chunk.
- Logically edited data items in previous chunks must be copied.
- All undeleted ancestors of logically edited or deleted data items must also be copied.

- Existing, unedited dependent data items of edited parent data items do not need to be copied. The new parent data item copy can reuse the positional data for the unedited dependent data items in the previous chunks.

20

- The chunk will contain data indicating the position of a new root data item.
- The chunk will contain data indicating the position of a previous chunk.

25

Thus each database transaction is embodied in a chunk of data, and the DBMS can directly reveal the entire state of the database as it was immediately after a transaction of interest by traversing a network of data items starting from a root data item indicated in the corresponding transaction's chunk of data. In this way the DBMS can directly reveal a snapshot of the database as it was at any moment in between transactions in the database's history.

30

In one embodiment of the invention the data values which occur frequently, or are likely to occur frequently in the logical database state, may be stored in data items in a part of the network of data items, in such a way that they can be reused as dependent

data items of many other data items. For example, textual values such as "London" may occur frequently in a database of UK addresses.

Furthermore, in an embodiment of the invention, the network of data items may 5 be in the form of a traditional network database. Generally a network database allows data records to be linked together in ways appropriate for application programs. Whilst traditional network databases conventionally do not enable old versions of the database to be directly revealed, it is possible to ensure that the DBMS of this embodiment enables all old versions of the network database to be revealed directly.

10

For example the network of data items may constitute a traditional relational database incorporating tables, views, columns, rows, fields, etc. The DBMS may contain an interpreter for standard query languages (SQL) which would enable users to access and modify a fully functional relational view of the database. In addition users 15 would be able to directly access a relational view of the database at all moments in between transactions in the database's history.

Furthermore the network of data items may constitute an object database. Generally an object database contains objects which consist of encapsulated data and 20 programmatic behaviour. In this case, the network would contain data items corresponding to object classes and object instances. Logical references between the objects would be modelled by the data items corresponding to the referring objects containing data indicating the positions of the data items corresponding to the referred objects.

25

In another embodiment of the invention, the network of data items may constitute a virtual disk drive (VDD), with the extra functionality of being able to directly access the logical state of the VDD at any time in its past. Generally a virtual disk driver is a computer software interface which allows other computer programs to 30 treat the software implementation of that interface as if it were a computer disk, so that such an interface typically supports functions such as reading, writing and modifying data at random positions within a certain range of valid positions. Preferably certain

data items in the network of data items would correspond to ranges of positions within the virtual disk (e.g. disk sectors). These data items would contain the data that was logically stored in the corresponding region of the virtual disk. If a computer application wished to modify the contents of a region in the VDD, then the DBMS 5 would prepare a new transaction chunk corresponding to that logical disk modification. The DMBS would prepare the chunk by locating the data item or data items corresponding to the region in the VDD, and creating new versions of those data items. Preferably the DBMS would sometimes split or merge the logical ranges corresponding to physical records to increase the efficiency of the physical data in representing the 10 corresponding logical data. Thus the amount of virtual disk embodied in different records may vary. Preferably the parent/dependent hierarchy of data items in the network of data items would relate to the locations of the corresponding regions and sub-regions on the virtual disk.

15 In a further embodiment of the invention a version control system (VCS) may be incorporated in the DBMS. Such a VCS would allow versions of the database be arranged in series to show how the database developed. The VCS can also contain branch points at which alternative versions of the logical state of the database are allowed to develop in parallel.

20 Preferably each new chunk of data contains data indicating the position of a previous chunk of data. This previous chunk of data may or may not be the chunk immediately preceding the new chunk in the file. The previous chunk may be an even earlier chunk. In this way, the VCS can arrange the chunks into a logical tree of 25 versions. Preferably each new chunk will also contain metadata such as the time it was created, the name of the user who made the change, the motivation of the user making the change, and the project, job or business associated with the change.

30 Certain data items in the network of data items, called version data items, may, along with their descendant data items, embody a particular version of a logical database. The version data items may themselves be dependent data items of version control data items. Preferably the VCS can navigate the network of data items from a

root data item, via the version control data items, to the version data items, and thence to the logical database data items.

Generally, in a multi-user, transactional database, several users (who may be 5 humans or other computer programs) can access the database simultaneously. If a user wishes to modify the database, the user begins a transaction, makes the necessary modifications, and then attempts to commit the transaction. If several users wish to modify the database simultaneously, then one or more of the users may have their modification request rejected by the DMBS (either at the begin stage, or the commit 10 stage).

In a development of the invention a multi-user DBMS is provided which avoids ever having to reject a modification due to several users requesting a modification simultaneously, by incorporating the VCS functionality related earlier into the 15 begin+commit transaction logic. Preferably, at the instant when a user begins a transaction, the DBMS can associate the chunk associated with that user's logical view of the database at that instant with the transaction. Then, when the user commits their transaction, the DBMS can append a new chunk which, when viewed via the VCS, logically follows on from the chunk associated with the transaction. If several users 20 have transactions open simultaneously, and they subsequently commit their transactions, then the DBMS may need to store the new versions in different branches. These branches can be reconciled later, possibly using application specific algorithms.

Generally application programs often provide users with undo/redo commands 25 permitting users to make changes to the database (e.g. a word-processing file) confident that they can undo any changes made, and then redo them if they change their mind. This frees users from some of the undesirable consequences of making mistakes.

In a further development of the invention VCS functionality is incorporated into 30 the DBMS to support the data management of an application that provides the user with an undo/redo mechanism. Preferably, as the user makes changes, the DBMS adds transactional chunks to the database. If the user uses the undo command, the VCS

within the DBMS preferably reverts to an earlier version of the database, so the user will see the modification being undone. Preferably, if the user then makes a different change, the DBMS appends another transactional chunk, and the VCS creates a new branch for that change. In this way all the database states which the user causes, and 5 the order of those states, are recorded. Thus the DBMS automatically collects the raw data required to analyse the behaviour and effectiveness of the user, and the mistakes made by the user. This raw data can be used to monitor the users' performance, help train users, and improve the user interface of the application software.

10 In an application of the invention, the DBMS may use append-only unmodifiable media to physically store the database, and yet present a logical view of the database which can be modified. For example, some types of compact disk can have data appended, but cannot modify data which has already been written. These types of append-only media are sometimes referred to as write-once-read-many 15 (WORM) devices.

In order that the invention may be more fully understood, reference will now be made, by way of example, to the accompanying drawings, in which:

20 Figure 1 shows schematically the physical structure of chunks in the database file of a DBMS in accordance with the invention;

Figure 2 shows schematically the logical and physical states of such a database during a transaction;

25 Figure 3 shows schematically how a DBMS in accordance with the invention can logically structure records to represent a general relational database;

Figure 4 shows schematically how a VCS can arrange chunks in a database file 30 of a DBMS in accordance with the invention; and

Figure 5 shows schematically how a VCS can logically structure records to

represent a general relational database with version control of a DBMS in accordance with the invention.

It should be understood that the figures are intended to show only the structure
5 of simple exemplary DBMS's in accordance with the invention by way of illustration of
the principles underlying such a system, and that actual systems which are likely to be
produced in accordance with the invention will incorporate additional levels of
structural complexity including additional features which would be well understood to
those skilled in the art.

10

Figure 1 shows how the basic structure of such a DBMS involving formatting a
database file as a series of chunks where a new transaction each day causes a new chunk
to be appended to the file. The boxes marked day1, day2 and day3 show the chunks
which are appended onto the file each day.

15

Figure 2 is an exemplary embodiment showing how the logical and physical
states of a database are related during a database-modifying transaction, in the case
where the database in question is a simple network database. The top part of the figure
shows the logical states, and the bottom part of the figure shows the corresponding
20 physical states. The lefthand side of the figure shows the state of the database on day 1
before the transaction, and the righthand side of the figure shows the state of the
database on day 2 after the transaction. Each transaction chunk preferably contains the
position of a root data item. The root data item for each chunk in each physical state
diagram is indicated by a black semicircle.

25

Before the database-modifying transaction, the database contains six data items
containing the names of six regions of the world: England, America, Africa, Canada,
Spain and France. In the logical state the data items are presented in a binary tree. In
this example, the binary tree is sorted, which means that every parent data item is
30 alphabetically later than all of its lefthand side descendants, and alphabetically earlier
than all of its righthand side descendants. In the physical state, the data items are stored
inside a single chunk. In the physical state, the parent data items (England, America

and Spain) will also contain data indicating the position of the dependent data items (America, Spain, Africa, Canada and France).

In this example, a user wishes to add a new data item "Turkey" into the 5 database. Accordingly the new data item "Turkey" is inserted into the new chunk. Since, in this example, the DBMS wishes the binary tree to remain sorted, the new Turkey data item will be inserted into the logical state as the righthand dependent data item of a Spain data item. This means that the old Spain data item must be copied, and the copied data item is labelled S* in the diagram. Similarly the old England data item 10 is copied as E*. Thus the new chunk will physically contain three data items: E*, S* and Turkey. The new E* data item will have America and S* as its dependent data items. The new S* data item will have France and Turkey as its dependent data items. The diagram shows how the two different logical states of the network database (i.e. before and after the transaction) can be directly revealed by traversing the network from 15 the root data item of one of the two chunks.

Figure 3 is a representation of a general purpose relational database as a network 20 of data items for use with a DBMS in accordance with the invention. In this example, there are different types of record, corresponding to traditional elements of relational database, such as strings, tables, rows, fields, column definitions and data values. Each table in the relational database has a corresponding table record. The table records are arranged into a sorted binary tree. In this exemplary scheme, each table record has up to 25 two dependent table records. Each table record also contains data indicating the name of the table within the relational database. This binary tree structure of table records does not have an analogous structure in classical relational database theory. In classical relational database theory, tables are considered to be more independent, with relationships between tables being inferred as-and-when required with join operations. The binary tree structure is used here to help the DBMS locate a table from its name.

30 Each table record also has a dependent record which forms the local root of a sub-network of row records. Each row record contains data which appears in a row in the relational database table corresponding to the table record in question. The diagram

only shows the row sub-network for one of the table records, although it should be understood that each table record preferably contains its own row sub-network. Similarly the diagram shows how each table record contains its own column definition sub-network, and each row record contains its own field sub-network, and each field record contains data indicating the field value.

5 In the example of Figure 3, there is a separate sub-network for strings. This sub-network contains canonical records for commonly occurring text values. Thus, if many relational database fields have the value "London", then all the corresponding field 10 records can store the position of a single, canonical string record representing "London".

Classically, a relational database is presented as a collection of tables. This tabular structure can be transformed into a network structure. There may be several ways to achieve this transformation, and one possible way is shown in Figure 3.

15 Whatever transformation is used, once a relational database or network database or object database or virtual disk drive or any other form of applicable database is converted into a network structure, the principles underlying present invention enable historical tracking and version control functions to be added.

20 Figure 4 shows an example of how the logical and physical states of a database are related during alternative simultaneous database-modifying transactions in a DBMS in accordance with the invention, for the case where the database is a simple network database. The top part of the figure shows the logical states, and the bottom part of the figure shows the corresponding physical states. The lefthand side of the figure shows 25 the state of the database on day 1 before the alternative simultaneous transactions, and the righthand side of the figure shows the possible states of the database on day 2 after these transactions. Furthermore each transaction chunk preferably contains the position of a root record. The root record for each chunk in each physical state diagram is indicated by a semicircle.

30 In this example the database contains two records on day 1 containing the names of two regions of the world: England and America. Furthermore two different users A

and B wish to make simultaneous additions to the database. User A wishes to add the record France, and user B wishes to add the record Germany. For the sake of this example, it is assumed that adding France and Germany are mutually exclusive options within any one logical database state.

5

The righthand side of Figure 4 shows at the top the two alternative logical databases which would result on day 2 from these additions. The bottom righthand side of Figure 4 shows how both additions can be physically logged in the database. This is done by adding two more chunks, corresponding to the two different transactions. Each 10 transaction is based on the initial chunk. According to the principles underlying the present invention each chunk contains data indicating the position of a previous chunk. As the diagram shows, both of the new chunks will contain data indicating that their previous chunk is the first chunk.

15

Figure 5 is an elaboration of the relational database schema shown in Figure 3 for use with a DBMS in accordance with the invention. In this example the network schema has new record types for version control records, and version records. These new record types allow the VCS to track historical versions of the relational database on different development branches, as well as backwards and forwards in linear 20 development steps.

25

CLAIMS:

1. A database management system for maintaining chunks of data indicative of the states of the database both before and after a transaction modifying the state of the database, the system comprising:
 - (a) memory means for holding permanent records of the state of the database before the database-modifying transaction and the state of the database after the database-modifying transaction;
 - (b) relation determination means for relating at least one parent data item in the record of each database state to at least one dependent data item in the record of the same database state;
 - (c) root determination means for determining the position of a root data item in the record of each database state to which other data items in that record are related; and
 - (d) state determination means for determining the state of the database after the database-modifying transaction by relating the root data item corresponding to that database state to both at least one data item in the record of that database state and at least one data item in the record of the state of the database before the data-modifying transaction as determined by the previous state location means.
- 20 2. A system according to claim 1, wherein the state determination means is arranged to relate the root data item to at least one dependent data item by way of at least one parent data item by use of the relation determination means associated with that parent data item.
- 25 3. A system according to claim 2, wherein the state determination means is arranged to record the position of the parent data item corresponding to each dependent data item during the tracking of data items.
4. A system according to claim 1, 2 or 3, wherein new record compiling means is provided to compile a supplementary chunk of data indicative of the state of the database after the database-modifying transaction and is arranged to copy those data items from the previous record which have been modified by the transaction whilst not

copying those data items from the previous record which have not been modified by the transaction.

5. 5. A system according to claim 4, wherein the new record compiling means is arranged to copy dependent data items from the previous record which have been modified by the transaction, as well as parent items to which those dependent data items are related by the relation determination means.
10. 6. A system according to any preceding claim, wherein presentation means is provided to present the data items in each record in a different logical structure.
7. 7. A system according to 6, wherein the presentation means is adapted to present the data items in the form of a relational database.
15. 8. A system according to 6, wherein the presentation means is adapted to present the data items in the form of an object database.
9. 9. A system according to 6, wherein the presentation means is adapted to present the data items in the form of a virtual disk drive.
20. 10. A system according to any preceding claim, wherein previous state location means is provided to relate the record of the state of the database after the database-modifying transaction to the position of the record of the state of the database before the database-modifying transaction.
25. 11. A system according to any preceding claim, which incorporates a version control system (VCS) defining branch points at which alternative versions of the logical state of the database are allowed to develop in parallel.
30. 12. A system according to claim 11, which is a multi-user system permitting several users to modify the database simultaneously to produce alternative versions of the state of the database after modification, wherein the memory means is adapted to

permanently hold a record of the state of the modified database produced by each user together with an indication of the user's logical view of the database before modification.

- 5 13. A system according to claim 11, which provides the user with an undo/redo mechanism, wherein the memory means is adapted to permanently hold records of the state of the modified database produced by first and second database-modifying transactions, whereby, after a first database-modifying transaction made by the user, the state of the database before such database-modifying transaction may be determined in
- 10 response to an undo command from the user, and subsequently the state of the database after a second database-modifying transaction different to the first database-modifying transaction may be determined in response to a redo command from the user, as an alternative to determination of the state of the database after the first database-modifying transaction.
- 15 14. A system according to 13, wherein analysing means is provided to analyse database-modifying transactions made by the user.
- 16 15. A system according to 13 or 14, wherein mistake identifying means is provided to identify common mistakes made by user in making database-modifying transactions.
- 20 16. A system according to any preceding claim, wherein each record contains metadata providing information relating to the creation of the record.
- 25 17. A database management system for maintaining chunks of data indicative of the states of the database both before and after a transaction modifying the state of the database, the system being substantially as hereinbefore described with reference to the accompanying drawings.
- 30 18. A programmed computer incorporating a database management system according to any preceding claim.

19. A data storage medium incorporating data recorded by a database management system according to any preceding claim.

ABSTRACT**“Database Management Systems”**

5 A database management system is adapted to hold permanent records of the states of a database both before and after database-modifying transactions so as to allow all previous states of the database to be directly accessed, without requiring storage of large amounts of redundant information. To this end the system relates parent data items (such as America, Spain; S*) in the record of each database state to dependent 10 data items (such as Africa, Canada, France; France, Turkey) in the record of the same database state. Additionally the system relates a root data item (such as England; E*) in the record of each database state to the other data items in that record. Such relationships allow the state of the database after a database-modifying transaction to be determined by relating the root data item (E*) corresponding to that database state to 15 both data items (S*, Turkey) in the record of that database state and data items (America, Africa, Canada, France) in the record of the state of the database before the data-modifying transaction.

(Figure 2)

1/4

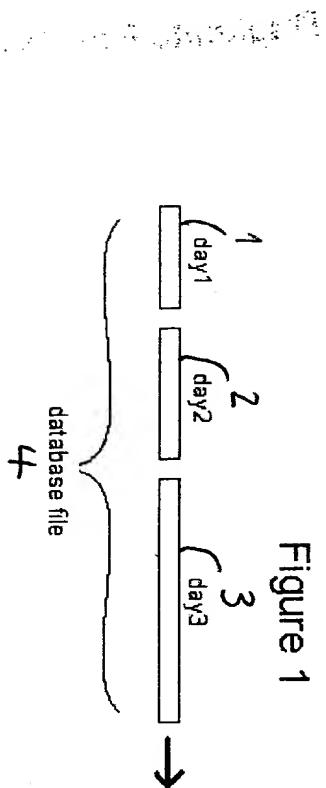
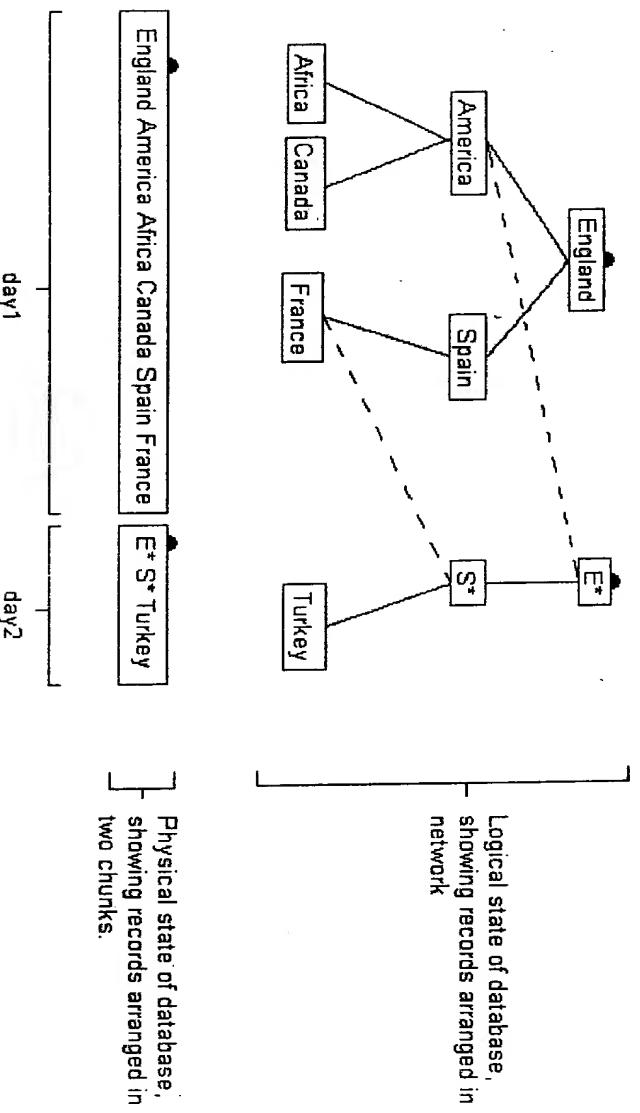


Figure 2



2/4

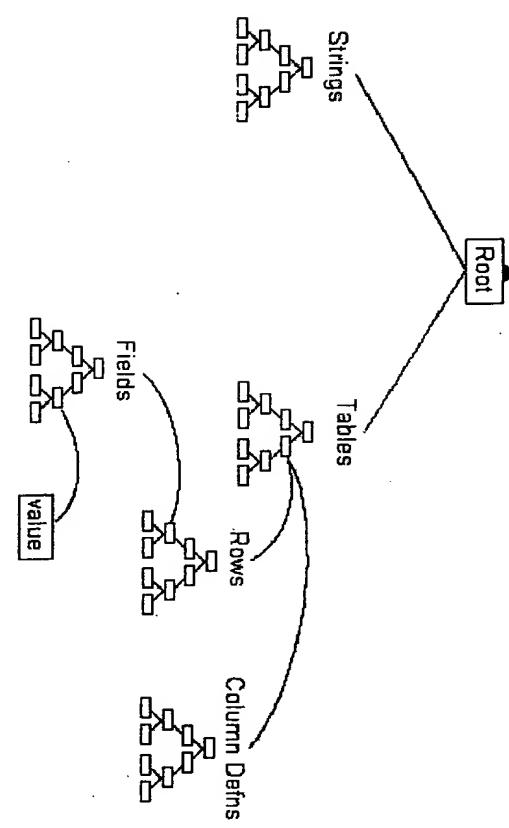
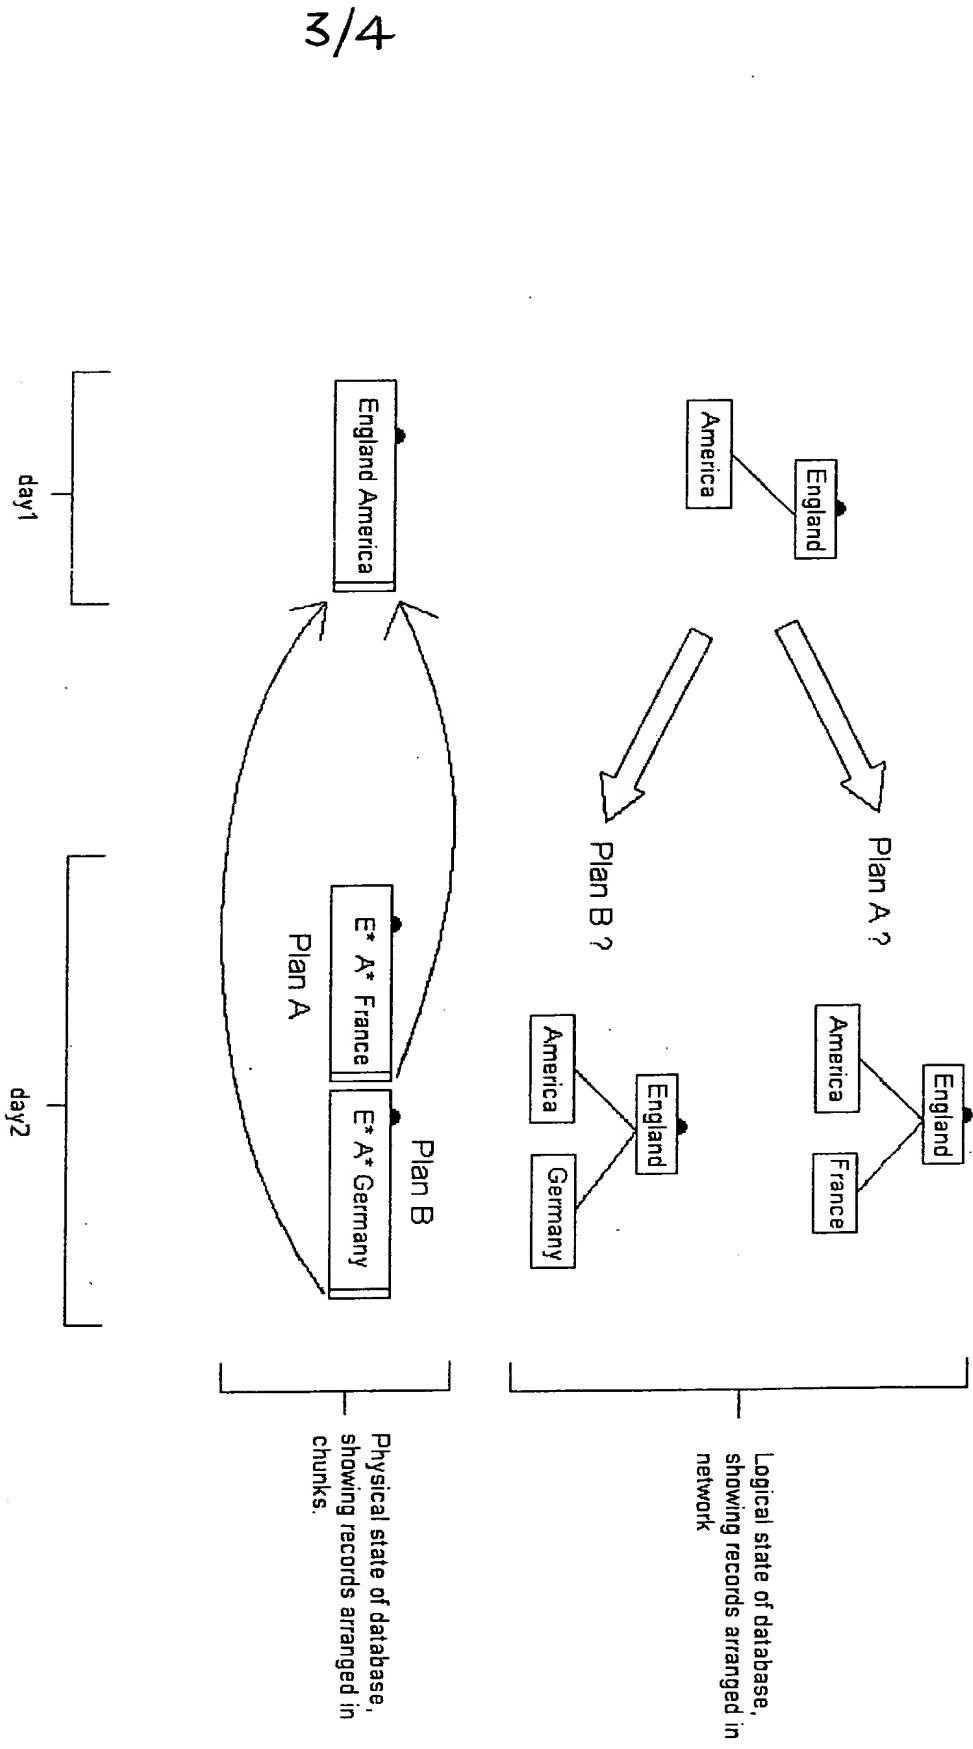


Figure 3

Figure 4



4/4

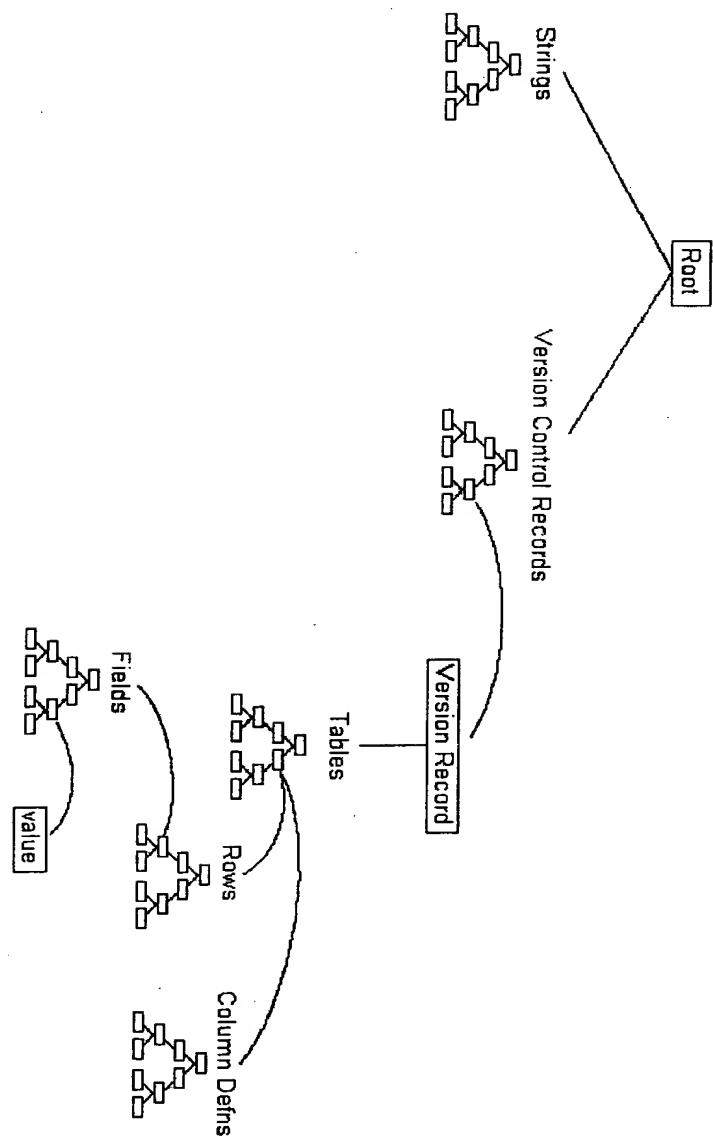


Figure 5